Decision Analytic Methods in RODOS

Vladimir BORZENKO
Institute of Control Sciences, Profsoyuznaya 65, Moscow, 117809 Russia
Simon FRENCH[1]
School of Computer Studies, University of Leeds, Leeds, LS2 9JT, UK

Abstract. In the event of a nuclear accident, RODOS seeks to provide decision support at all levels ranging from the largely descriptive to providing a detailed evaluation of the benefits and disadvantages of various countermeasure strategies and ranking them according to the societal preferences as perceived by the decision makers. To achieve this, it must draw upon several decision analytic methods and bring them together in a coherent manner so that the guidance offered to decision makers is consistent from one stage of an accident to the next. The methods used draw upon multi-attribute value and utility theories.
Keywords: Constraint satisfaction; decision support systems; expert systems; HERESY; M-Crit; multi-attribute value and utility theory.

1 Introduction

The lack of a uniform response to the Chernobyl accident, both in and beyond the former Soviet Union, has led to a number of projects supported by the Commission of the European Communities (CEC), under its Radiation Protection Programme. RODOS (Real-time On-line DecisiOn Support system) is one of these projects. It is designed to be a comprehensive decision support system (DSS) for off-site emergency management, which will provide support from the moment that an accident threatens through to long term countermeasures implemented months and years after an accident. A key feature of RODOS is that it seeks to provide decision support at all levels ranging from largely descriptive reports to a detailed evaluation of the benefits and disadvantages of various countermeasure strategies and their ranking according to the societal preferences as perceived by the decision makers: see Table 1. To provide such comprehensive decision support many design issues need to be addressed. In this paper, we describe the architecture of RODOS with specific reference to modules in the evaluation subsystems (ESY). For a more general introduction to the design of RODOS, we refer to papers earlier in this session, specifically [1], [2], and to [3].

Level 3 decision support (Table 1) requires complex modelling of preferences and values. The design of RODOS uses multi-attribute value and utility functions (MAV/UT) to provide this: see [4], [5] and [6] for discussions of these methods. Two modules, HERESY and M-Crit have been developed. These implement MAV/UT methods in subtly different

[1] Please address correspondence to Simon French

Table 1 Levels of decision support for off-site emergency management. Decision support can be provided at various levels, here categorised into four levels. The functions provided at any level include those at lower levels. RODOS is unique in that it will provide support at all levels, including Level 3 for all potentially useful countermeasures at all times following an accident.

Level 0:	Acquisition and checking of radiological data and their presentation, directly or with minimal analysis, to decision makers, along with geographic and demographic information available in a geographic information system.
Level 1:	Analysis and prediction of the current and future radiological situation (i.e. the distribution over space and time in the absence of countermeasures) based upon monitoring and meteorological data and models.
Level 2:	Simulation of potential countermeasures (e.g. sheltering, evacuation, issue of iodine tablets, food bans, and relocation), in particular determination of their feasibility and quantification of their benefits and disadvantages.
Level 3:	Evaluation and ranking of alternative countermeasure strategies in the face of uncertainty by balancing their respective benefits and disadvantages (e.g. costs, averted dose, stress reduction, social and political acceptability) taking account of societal preferences as perceived by decision makers.

ways. The design of RODOS also makes use of constraint management and expert system technologies within the ESY subsystems to help in problem structuring and in explaining to the decision makers the guidance provided by M-Crit and HERESY modules in formulating and evaluating countermeasure strategies.

The organisation of this paper is as follows. We begin by discussing a little further the decision support required during a nuclear accident. Section 3 illustrates multi-attribute modelling of consequences. Section 4 describes the ESY subsystem and the modules which form it, particularly the two MAV/UT modules: M-Crit and HERESY. Finally, the concluding section notes several issues relating to decision support, particularly matters related to the validation.

2 Decision support during a nuclear accident

RODOS is designed to support decision makers throughout all phases of a nuclear accident. Initially, RODOS will support the decision making of plant or site managers and local emergency management. Later, regional or national governments will become responsible for decision making, depending on how severe the accident is. Thus RODOS will support all decision makers and all decision making on countermeasures from initial evacuation, sheltering and issue of iodine tablets, through food bans to long term relocation.

The need to provide such comprehensive support introduces a number of issues that are seldom faced in the design of DSS's for other contexts.

- *Multiplicity of decision makers.* Many decision makers will be involved in the emergency management: plant managers, the emergency services, regional emergency planning officials, local, regional and national politicians. Each has differing levels of technical competence and differing information needs. Perhaps more importantly, they will have differing levels of authority to express value judgements. Plant managers will not be able to 'speak for the public', whereas politicians do have that authority.

- *Multi-criteria, public equity and risk.* There is no single criterion for choosing between countermeasure strategies. In addition to those directly relating to health risks arising from the radiation, there are issues related to psychological stress, public acceptability, and equity of risk sharing across the population. Equity issues are discussed in [8].
- *Many different levels of urgency.* Initially, decisions on countermeasures must be made in a matter of minutes; later, the timescales 'relax' with decision making able to take several hours or days.

In the early phases of an accident, the urgency and the probable lack of involvement of political decision makers means that issues related to 'intangibles' such as public acceptability and equity of treatment must be pre-programmed in some sense. Moreover, during this phase health issues related to the stochastic and non-stochastic effects of potential exposures will drive the decision making. In the later phases, decision making will have to address political and social imperatives in addition to pure health issues. This means that the support offered by RODOS must vary both in the range of criteria used in the analysis and in the software interface. Moreover, looking at Table 1, we may expect that full level 3 support may not be used in the first few hours of an accident.

3 Multi-attribute modelling of consequences

We shall use the term *attribute* to mean one of the dimensions along which we assess the consequences of a decision. Thus `collective dose`, `individual dose` and `cost` are attributes. A *criterion* or *objective* refers to an attribute and a direction of preference. Thus `minimise collective dose`, `minimise individual dose` and `minimise cost` are criteria. Some attributes are objective in that they correspond to physical measurements; others are more subjective requiring judgement in their definition, e.g. equity of treatment. It is conventional to organise the attributes involved in analysis into hierarchies. This offers many cognitive advantages and also helps structure decision analyses [6].

Figure 1 gives two attribute hierarchies that have been developed in case studies on relocation decisions after an accident. Note that some of the attributes are clearly objective, whereas others must inevitably reflect subjective judgements. Discussions of how these were defined and used to structure sensitivity analyses may be found in [4] and the

Normal Living
- Resources
- Effects
  - Health
    - Radiation Related
      - Fatal Cancers
      - Hereditary
    - Stress Related
  - Public Acceptability
    - Affected Region
    - Rest of USSR

Hierarchy from the Chernobyl Study

Overall
- Cost
- Effects
  - Health
    - Collective Dose
    - Individual Dose
    - Psychological
  - Social/Political Acceptability
  - Flexibility

Hierarchy from the BER-3 case study

Figure 1: Attribute hierarchies developed in two multi-criteria decision analyses on countermeasure strategies [4]

references cited therein. The use of event conditional attribute hierarchies to support decision making when there is a threat of an accident is discussed in [8].

4 The ESY subsystem

The conceptual architecture (not the physical software which includes, e.g., databases and user interfaces) of RODOS consists of three types of subsystem [1], [2], [3]:

- ASY (analysing subsystem) modules process incoming data and forecast the location and quantity of contamination including temporal variation.
- CSY (countermeasure subsystem) modules suggest possible countermeasures, check them for feasibility, and calculate their expected benefit in terms of a number of attributes (criteria).
- ESY (evaluation subsystem) modules rank countermeasure strategies according to their potential benefit and preference weights provided by the decision makers.

We focus on the ESY subsystems of modules, which support the evaluation of different countermeasure strategies. These implement the level 3 support offered by RODOS. Rules, weights and preference functions are encoded in the ESY and applied to a list of alternative countermeasures to provide a ranked short list to decision makers. Both the ASY and the CSY will use several models throughout the accident depending on the time, the location and the actual situation. However, the ESY may be based upon the same software module with different attribute trees and with the preference weights changing over time.

An ESY subsystem will operate in interactive mode using graphical interfaces to communicate with a variety of decision makers who may possess many qualitatively different skills and perspectives: e.g. scientists, medical personnel, engineers, emergency planners, government officials and senior politicians. It will present the countermeasures in a ranked short list together with those rules and preferences that determined the order of the list. Intuitive justifications for choices and underlying uncertainties inherent in the predictions will also be provided. The ESY will assist users in modifying rules, weights and preferences and other model parameters as well as exploring the consequences of each change. The importance of this exploration cannot be overemphasised. Any DSS helps decision makers not by *making* the decision itself, but by *enhancing the decision makers' understanding* of the problem, the issues before them and their value judgements. They are then better able to make the decision because of this greater understanding.

4.1 Architecture of an ESY subsystem

The ESY will[2] have the form of Figure 2. It comprises three further subsystems:

- A coarse expert system filter which rejects any strategies that are logically infeasible or do not satisfy some given constraints.
- A multi-attribute utility theory (MAV/UT) ranking module which takes the remaining list of strategies as input. It ranks the strategies for their relative effectiveness according to previously elicited utility attributes and preference weights from the decision makers. It may be necessary to revise and re-evaluate these preference weights in any given situation before a particular decision is taken.

---

[2] Prototypes of the ESY modules are written and are currently being evaluated; however, they are not fully implemented into the current realease of RODOS.

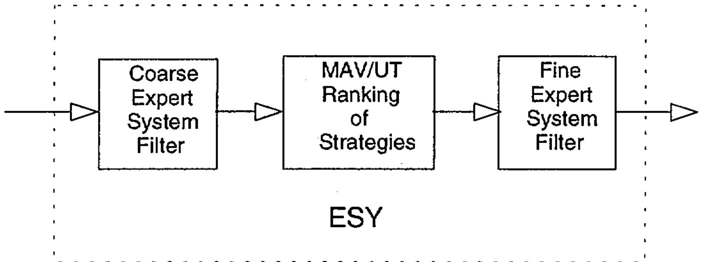

Figure 2 The conceptual structure of the ESY module

- A fine expert system filter which takes the top 10-15 strategies and produces a management summary report detailing the costs and benefits of each.

4.2 Constraint management coarse expert system

Although the number of protective measures which may be taken is limited, the number of potential countermeasure strategies is enormous. This arises because the region will be divided into a number of areas and a strategy specifies which measures should be applied in each of these areas. Moreover, each measure may have a time period associated with it, e.g. start and finish times for sheltering. Thus the number of strategies can grow combinatorially. However, very few of these conceptually possible strategies will be worthy of close examination. They may run break the guidance provided by intervention levels or may be infeasible in practical terms (e.g. requiring 100,000 people to be evacuated in 30 mins). Moreover, they may break simple principles describing public acceptability, e.g. protecting children less effectively than adults or evacuating lower risk areas in preference to higher ones. One would never evacuate part of a small village; the public would never accept or understand such an action. There must be continuity of treatment.

A very simple expert system will be used to discard strategies which are incompatible with such principles. This coarse expert system is being implemented using constraint management techniques [9], [10]. This technology is, in one sense, as old as combinatorial programming, for it does nothing other than identify objects that satisfy a set of constraints. But, in another sense, it is very new in drawing upon modern tree search and list manipulation algorithms implemented with artificial intelligence languages. Early experiments show that constraint satisfaction technology is well able to cope with the combinatorial problems we face here. In a simple example with a potential 17 billion possible countermeasure strategies were reduced to a list of about 500 which were worthy of further evaluation. The reduction took a matter of seconds on a Sun Sparcstation of comparable power to the HP workstation on which the full RODOS system runs.

The strategies satisfying the constraints imposed by the coarse expert system will be passed to a MAV/UT ranking module, which will identify the top ten or twenty ranking strategies. Two such modules have been written: M-Crit and HERESY.

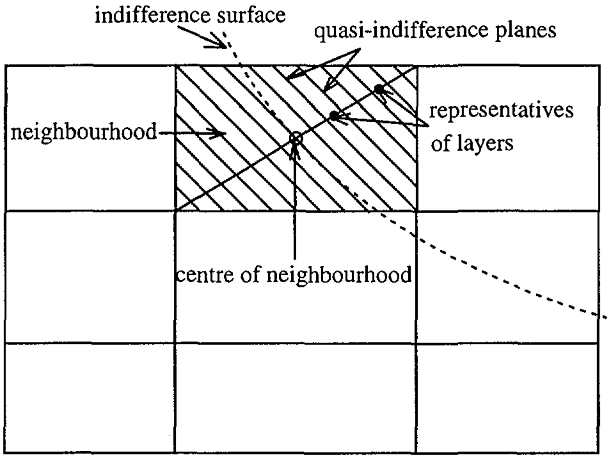

Figure 3 Splitting of the region of feasible solutions into neighbourhoods and the approximation of indifference planes by linear quasi-indifference planes within a neighbourhood

4.3 M-Crit

M-Crit implements a piecewise linear approximation to the decision makers' indifference planes: an indifference plane is a surface or contour of points of equal value to the decision makers. The method is know by the acronym PLANT (piecewise linear approximation numbering technique) [11]. It assumes that the decision makers have a well formed set of preferences and that the problem is to model these and articulate them in the context of a particular problem. It approximates the decision makers' indifference surfaces by splitting the area of feasible solutions in the criteria space into a number of rectangular neighbourhoods; see Figure 3. Within each neighbourhood, the indifference surfaces are approximated linearly by eliciting the substitution coefficients from the decision makers. By 'chaining together' the approximating linear indifference (hyperplanes) the method can approximate complex preference structures to a reasonable degree of accuracy.

M-Crit is a window-based interactive implementation of the PLANT method. It allows the approximation to be constructed interactively and displays the approximation back to the decision maker visually for checking. While this visual reflection does allow the decision makers to consider whether their preferences are appropriate, no other consistency checking is built into the elicitation procedure. Nor are there any underlying preferential independence assumptions such as additive or utility independence which are often required in MAV/UT modelling [5], [6]. This makes the method appropriate for circumstances in which decision makers are sure of their value judgements.

4.4 HERESY

An alternative approach is taken by the HERESY module. This implements MAV/UT models based upon much stronger assumptions: e.g. preferential independence. These assumptions

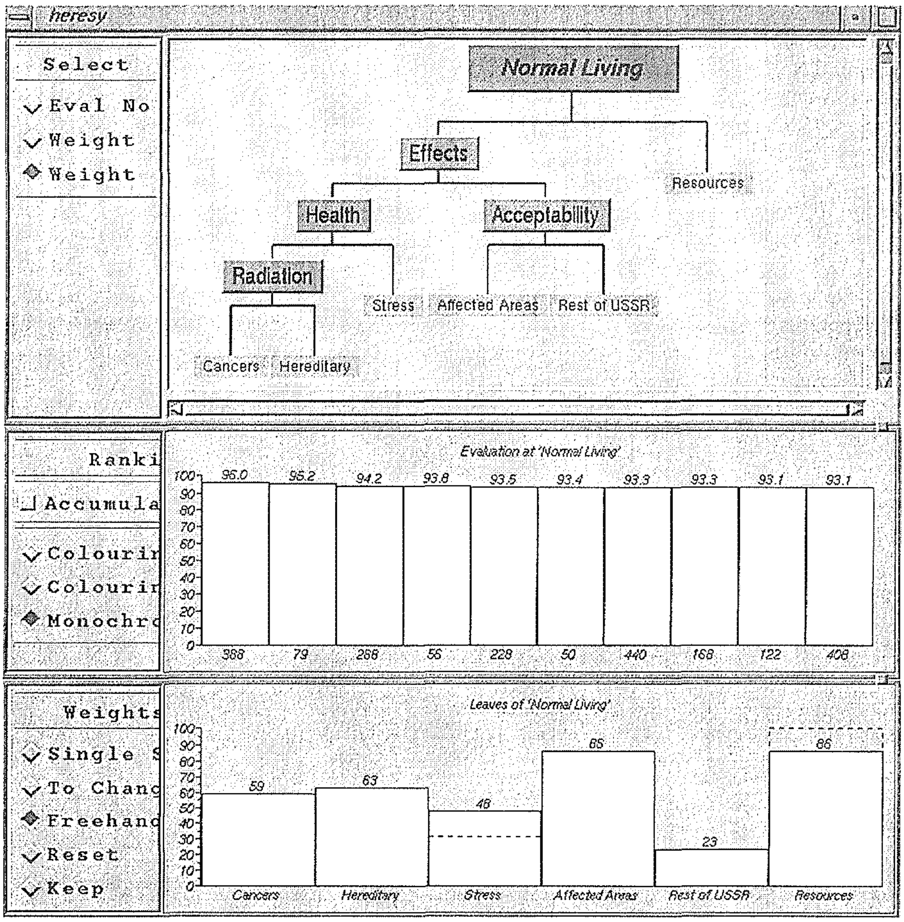

Figure 4 Screen dump of the HERESY module

restrict the form of value and utility functions considerably. The current implementation assumes sufficient independence to ensure additivity[3], but future implementations will also allow multiplicative and multi-linear forms. The advantage of introducing such restrictions is that the elicitation procedures can include more consistency checking and more detailed sensitivity analyses may be performed after the evaluation. The disadvantage is that the decision makers may be unable to articulate some of the value judgements which they might wish to express. Thus the strengths and weaknesses of M-Crit and HERESY complement each other.

HERESY's purpose is to identify the top few, say ten, ranking strategies and check the sensitivity of these to the choice of weights on different criteria. A screen dump is shown in Figure 4. The screen is divided into three areas. At the top the attribute hierarchy is shown.

[3] In an additive MAV model the overall value or score of an alternative is formed as a weighted sum of scores on individual attributes.

In the middle is a histogram showing the overall scores (values) of the top ten strategies. At the bottom is a histogram showing the current weights on (some of) the attributes. Not all the weights need be shown simultaneously. There may be cognitive advantages in concentrating attention on particular branches within the attribute hierarchy. All bars on the histogram are labelled appropriately. The user selects a weight with a mouse by clicking on the appropriate bar in the bottom histogram, and then increases or decreases the weight either by the keyboard or by pulling with a mouse. As the weight is changed, the middle histogram changes accordingly. When a change in the ranking of the top ten strategies occurs (or when one drops out of the top ten and another enters), the histogram rearranges itself. There is an audible beep and the user is informed of the change in a text window. Thus the user can identify the sensitivity of the ranking to the default weights in the model. The computational speed of the prototype confirms that the identification of the top 50 ranking strategies of about 10000 countermeasure strategies and associated sensitivity analysis can be performed almost instantly.

4.5 Fine expert system

After potential countermeasure strategies have been ranked using M-Crit or HERESY, the list of top ranking ones will be passed to an expert system with a sophisticated set of rules, each of which will be applied to each of the candidate strategies. The small number of strategies would allow a full set of explanations to be developed, which would give a critique of each of the strategies. Thus the output of RODOS will be a shortlist of strategies, each of which satisfies the constraints implied by intervention levels, practicability, etc., together with a detailed commentary on each strategy explaining its strengths and weakness. Klein [12] discusses a similar combination of expert system technology with MAV/UT ideas to provide decision makers with explanatory remarks on the ranking of strategies.

5 Discussion

The decision analytic issues involved in the design of RODOS are complex and we have only been able to touch upon a few of them here. A major omission is our lack of discussion of the relationship between uncertainty handling and preference modelling. Some discussion of this may be found in [7]; see also [13]. We have focused upon the prototype systems M-Crit and HERESY, which in their current implementations ignore uncertainty and risks. A later paper will report on the enhancement of their functionality to deal with these.

A more significant omission is the lack of discussion of validation and quality assurance. Validating a decision support system which one hopes and, indeed, plans should *never* be used brings a host of problems and issues. Moreover, the geographic and cultural spread of the many institutes involved in the software development itself raises many quality assurance issues. The latter issue of software quality assurance is being addressed very fully in the current fourth Framework Research and Development Programme of the CEC, during which quality assurance procedures are being applied. However, validation of the decision analytic methodologies is a more difficult matter. Essentially, we need to work very closely with a large number of decision makers to ensure that we are supporting them in the ways that they need.

Several exercises have been run and more are planned to explore and validate the use of RODOS in general and the ESY modules in particular. Our belief is that by working with

decision makers as they make decisions – in albeit artificial circumstances – we shall discover where the design and implementation of RODOS is poor or deficient. Already it seems to be clear that the decision makers do not perceive a need for level 3 support in the early phase of an accident [14], whereas they do find level 3 MAU/VT support useful in the decision making on later countermeasures such as relocation [4]. However, it is also clear that whilst the team designing RODOS must learn the needs of decision makers, the decision makers themselves need to learn the potential of modern DSS's. Currently, they are unaware of what is possible. None the less, it does seem that for the present the methods and software described above will be more useful after the immediate emergency has passed. For further discussion of the contribution that these methods may make, see [4].

6 Acknowledgements

The development of the MAV/UT methods for RODOS described here was supported by grants from the Commission of the European Community (B17-0060-GB, F13P-CT92-0036, F13P-CT92-0136, Sub94-F15-028). The ideas in this paper were developed in close co-operation with R. Borodin, A. Despres, J.E. Ehrhardt, N. Papamichail, D.C. Ranyard, L. Simpson, J.Q. Smith, E. Trakhtengerts and D. Vanderpooten. N Papamichail has developed a prototype of the coarse expert system using the ILOG constraint satisfaction toolkit.. S. Young programmed the prototype of the HERESY module and R. Borodin programmed the M-Crit module.

References

[1] J. Ehrhardt, V. Shershakov, M. Zheleznìak, and A. Mikhalevitch, RODOS: Decision support for off-site emergency management in Europe, First International Conference of the European Commission, Belarus, Russian Federation, and Ukraine on the Radiological Consequences of the Chernobyl Accident.

[2] O. Schüle and V Kossykh, The software environment of RODOS, First International Conference of the European Commission, Belarus, Russian Federation, and Ukraine on the Radiological Consequences of the Chernobyl Accident.

[3] J. Ehrhardt, J Päsler-Sauer, O. Schüle, G. Benz, M. Rafat and J. Richter, Development of RODOS, a Comprehensive Decision Support System for Nuclear Emergencies in Europe - an Overview, *Radiation Protection Dosimetry,* **50** (1993), 195-.

[4] S. French, Multi-attribute decision analysis in the event of a nuclear accident, *Journal of Multi-Criteria Decision Analysis* **5** (1996) in press.

[5] S. French, Decision Theory: an Introduction to the Mathematics of Rationality, ISBN 0-85312-682-8. Ellis Horwood, Chichester, 1986.

[6] R.L. Keeney. Value Focused Thinking: a Path to Creative Decisionmaking. ISBN 0-674-93197-1, Harvard University Press, Cambridge, Mass., 1992.

[7] S. French, D. Ranyard and J.Q. Smith, Uncertainty in RODOS, RODOS(B)-RP(94)05, Research Report 95.10, School of Computer Studies, University of Leeds. Available by connecting to WWW at file://agora.leeds.ac.uk/scs/doc/reports/1995 or by anonymous ftp from agora.leeds.ac.uk of the file scs/doc/reports/1995/95_10.ps.Z. 1995.

[8] S.French, E. Halls and D.C. Ranyard, Equity and MCDA in the event of a Nuclear Accident. ". RODOS(B)-RP(95)03. Contributed paper at the XIIth International Conference on MCDM at Hagen, June 1995. Accepted for publication in the Proceedings.

[9] P. van Hentenryck, Constraint Satisfaction in Logic Programming. MIT Press, 1989.

[10] E. Tsang, Foundations of Constraint Satisfaction. Academic Press, New York, 1993.

[11] V. Borzenko, E. Trakhtengerts, R. Borodin and V.M. Shershakov, M-Crit: a Multi-Criterial Decision Support Subsystem for RODOS Prototype 2. Scientific Production Association "Typhoon", Emergency Center, Obninsk, Russia.
[12] D.A. Klein, Decision-Analytic Intelligent Systems: Automated Explanation and Knowledge Acquisition. New Jersey: Lawrence Erlbaum Associates, 1994.
[13] V.M. Shershakov and E. Trakhtengerts, Development of RODOS/RECASS System as a Distributed Decision Making Support in Emergency. Institut of Control Sciences, Moscow, Russia. 1995.
[14] M. Ahlbrecht, J.E. Ehrhardt and S. French, Designing the Evaluation Module in RODOS/RESY: Execution and Analysis of Elicitation Exercises with Emergency Management Teams. School of Computer Studies, University of Leeds, Leeds.